



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

The U.S. Government Role in HPC: Technology, Mission, and Policy

J. R. Neely

April 17, 2014

Conference on Comparing High Performance Computing in
the U.S. and China
San Diego, CA, United States
April 29, 2014 through April 30, 2014

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

THE US GOVERNMENT ROLE IN HIGH PERFORMANCE COMPUTING: MISSION AND POLICY

*J. Robert Neely
Lawrence Livermore National Laboratory
neely4@llnl.gov*

*Prepared for the:
Conference on Comparing High Performance Computing in the U.S. and China
University of California Institute on Global Conflict and Cooperation
University of California at San Diego
April 29-30, 2014*

Abstract – The U.S. Government has played a critical role in defining U.S. leadership in HPC through the seven-decade history of computing. In the early years of computing, government spurred innovation in HPC technology largely through defense mission needs at the national laboratories. As computing emerged as a mass-market technology in the 1980's, the U.S. Government began formulating HPC policy in reaction to increasing global competition. Through a combination of initiatives and funding, the U.S. has maintained a global leadership role, but one that is increasingly tenuous as other countries are aggressively pursuing HPC as a competitiveness stimulant. This paper illustrates the history of supercomputing in the U.S. from the post-WWII to the present, outlines the role government has played, describes the agencies who have been paramount in executing mission through HPC, discusses the advisory groups influencing policy, and concludes with a look at what needs to be done today to continue exercising HPC leadership in the U.S.

INTRODUCTION

Since the very inception of computing, the United States has been a global leader in the development, deployment, and effective use of High Performance Computing (HPC). During much of this history this lead has been undisputed, at some times shaky, and most recently one that is not taken for granted. A number of factors have kept the U.S. at or near the top, including but not limited to: early and sustained leadership in computing across the usage spectrum, a strong set of mission drivers, a robust U.S. vendor community, an outstanding higher educational system, and strong support from federal agencies helping to drive adoption. But many worry that the lead we have held for so long could easily slip away from us without a coherent policy and funding roadmap for HPC. The fact that other nations have awoken to the benefits that HPC brings and are aggressively pursuing advanced computing

through coordinated policy and public-private partnerships adds to the urgency.

In the U.S., the importance of HPC leadership is largely viewed through the lens of the benefits that HPC provides, versus the goal of simply deploying a system deemed “the fastest”. Much like how building the world's fastest manned rocket ship may be a technological feat in and of itself, it would be impractical and wasteful to pursue unless it can take us somewhere useful or provide significant benefits to society that would not have happened otherwise. Likewise, extreme scale computers are a tool - a means to an end and not the end in and of themselves. But to the extent that they are used to discover new science, innovate new products, increase national security, and drive economic benefit, it is a tool that we as a nation want to maintain leadership in developing, as the primary innovator of HPC technologies will be first to benefit in the areas HPC directly impacts.

U.S. government has directly and indirectly impacted the technical direction of HPC primarily through two means. First, government has been a dominant investor both in procurement of large systems and in much of the up-front research necessary to field first-of-a-kind systems. These large procurements involve shared risk with the vendors, and ideally results in a broadly marketable product line within a year or two of initial deployment to the government. Most recently, this has been realized in the Cray XT and IBM BlueGene lines of Supercomputers – both of which started as first-of-a-kind machines for DOE laboratories, and were later successfully marketed to a broader audience by the respective vendors. Second, the government can issue broad policy decisions that provide both the focus and the funding guidance for addressing challenges of HPC. The latter is something that has largely emerged within the past 30 years once it became clear that the U.S. was not the only country pursuing HPC, and a government role in setting broad inter-agency directions was necessary to help provide a unified path forward with maximum leverage.

Is the U.S. really in danger of losing our global leadership role in HPC? Have we already? And how does one measure this? If we are in danger and want to ensure our leadership continues, how must government play a principal role? Many in the U.S. HPC community lament that the U.S. lacks a coherent HPC policy, and that without one – we will one day soon wake up and realize that we have lost our lead, including the benefits that come with that. As HPC increasingly pervades a larger and broader segment of our society, leadership has both direct and indirect consequences on a number of dependent domains, not the least of which is the ability to attract the best and brightest computer and data scientists in the world to choose to work in the U.S. for the betterment of U.S. interests.

This paper explores a largely historical perspective on how Government has influenced HPC in the U.S., its lasting effects, and where it is heading at the time of this writing. We first outline the history of supercomputing platforms in the U.S. and disruptive architectural changes that have occurred through the decades. The metrics (or lack of) for defining HPC leadership are then discussed, followed by an overview of how policy is executed in the U.S. government, the role of various advisory groups, and the major cabinet-level departments assigned to effectuate those policies. We then present a history of major U.S. policies and programs that have occurred, and their impact on boosting and maintaining U.S. HPC leadership. Finally, we present a set of high level policy recommendations apropos to the current global landscape in HPC.

A HISTORY OF SUPERCOMPUTER TECHNOLOGY IN THE U.S.

Before delving into policy issues, we will first take a brief tour through the history of HPC technology in the United States, and a sampling of the systems that defined the multiple eras of Supercomputing. This growth in the power of computers (a 50 billion-fold increase over approximately 70 years) is virtually unprecedented compared to other technologies, and has resulted in once unimaginable computing power available in something as simple as a smart phone, or even the embedded computers that help your automobile run efficiently.

WW II through the early 1960's

As with many technologies, computing in the U.S. was initially funded and developed out of military need, and thus driven by government. World War II spurred the development of the very first computers used to calculate ballistics projectiles. In fact, the word “computer” initially referred to humans (almost always women) serving the war effort by hand-calculating these tables. In order to

speed this process, the U.S. Army contracted development of ENIAC at the University of Pennsylvania’s *Moore School of Electrical Engineering*. Finished shortly after the conclusion of WWII, it was immediately put to use to help determine the feasibility of the Hydrogen Bomb by famed Princeton mathematician John von Neumann, a key figure in the Manhattan Project then consulting for Los Alamos. Von Neumann was one of the fathers of the theory of computing and developed the model for computer design that is still in use today, of which a CPU consisting of control and arithmetic logic is separate from the memory bank that holds both instructions and data feeding the CPU. His architecture was fully realized in the EDVAC computer, constructed shortly after ENIAC, and laid the foundations for the design of computers to this day. Los Alamos continued to drive the development of computing throughout the 1950’s with the development of ENIAC-inspired designs, including the *MANIAC I* and *II*.

While computing was taking root at Los Alamos, its sister laboratory Lawrence Livermore (then called the *Lawrence Radiation Laboratory*) was established in 1952 and procured a *Univac I* as one of its first purchases. Throughout the remainder of the 1950’s, Los Alamos and Livermore continued to purchase ever more-powerful computers, primarily from IBM (the 701, 704, 709, 7090). Each new computer brought some innovation in computer design, such as better memory technology or the inclusion of floating point arithmetic, culminating in the delivery of the first transistor-based machines – the *IBM 7090* and *Univac Livermore Advanced Research Computer* toward the end of the decade. The LARC was particularly interesting in that lab scientists contributed significantly to the design of the machine alongside Remington Rand’s engineers – one of the earliest examples of co-design. LANL’s response to the LARC was to build the *IBM 7030 “Stretch”*, which at the time was deemed a failure because it only achieved 50% of its aggressive performance goals. But many of the innovations developed in the Stretch system went on to form the foundations for the immensely successful *IBM 360* series.

The introduction of close collaborations between government labs and computing vendors who were having commercial success in academia and business, eventually led to what became a model for public-private partnerships for the design of leading-edge computers that lasts until this day.

The 1950’s and early 1960’s closed out an era of computing where the industry was really just beginning to find its way. Whether you desired a computer to do high-end scientific computing, or run your business – the mainframe computers of that era were all that was available. It wasn’t until computer designers began specializing in Supercomputers due to the demands primarily from the defense missions at the U.S. national

laboratories that the concept of High Performance Computing broke away from mainstream as its own category.



Figure 1 - The UNIVAC LARC (Livermore Advanced Research Computer) circa 1960. Developed by Remington Rand with design input from LLNL. (Picture courtesy of LLNL Archives)

The Introduction of the Supercomputer

No single person personifies the field of HPC in history as does Seymour Cray. His influence on computing from the late 60's through the late 80's defined the field in ways that has had lasting impacts on computing. Seymour Cray found his way early in the computing era working for *Control Data Corporation* (CDC) through his initial design of the *CDC 1604* in the early 1960's. The 1604 was amongst the first commercially successful transistor-based computers (along with the IBM 7090), and provided CDC with the momentum to compete head-to-head with IBM in the emerging scientific computing marketplace.

The first *CDC 6600* was delivered to CERN in Switzerland in 1965, and soon after to LLNL in 1966. The 6600 has since been dubbed the first “real Supercomputer” for its ability to outperform the IBM Stretch by a factor of about 3x, and its focus on performance through simplifying the tasks performed by the main CPU. Specifically, the 6600 separated input/output from the main task of computation, an innovation that ultimately led to the concept of *Reduced Instruction Set Computers* (RISC) architectures, which dominated HPC designs decades afterwards.

The *CDC 7600* was another Seymour Cray design that followed and built upon the success of the 6600 – and proved to be the fastest computer in the world for an astounding six year run from 1969 to 1975. In addition to incorporating faster clock speeds, the 7600 innovated the concept of an instruction pipeline, which allowed for low-level parallelism to be achieved. It also introduced the concept of multi-level memory, with a smaller “fast”

memory effectively acting as what would later be known as a cache.

The Vector Era (1970's – 1980's)

CDC continued to design and manufacture new models, most notably the *CDC Star 100*, one of the first computers to introduce the concept of *vector computing*. Vector computing was inspired by the success of high level programming languages that were becoming dominant in scientific computing, where *arrays* of data containing different values computed by the same set of instructions cried out for the efficiencies offered by the vector registers at the core of these designs. In vector computing, the similarity of arithmetic operations performed on arrays of data was highly optimized, allowing for efficient floating-point computations performed nearly concurrently in a dedicated instruction pipeline fed from high-speed vector registers. Despite the promises of vectorization, the disparity in performance between coding that could be successfully vectorized versus non-vectorized (or serial) code on the CDC Star 100 was so great that programmers found it difficult to achieve anything close to the promised performance.

By the mid-1970's, Seymour Cray had left CDC to start his own company, *Cray Research*, with a focus on HPC. Cray Research, thanks to a combination of engineering talent hired from nearby CDC and some financial backing from investors, introduced the *Cray-1* in 1975, with delivery to Los Alamos in 1976 and Lawrence Livermore in 1978. Unlike the CDC Star 100, applications ported to the Cray-1 saw an immediate performance gain, largely due to the attention paid to optimizing serial (non-vectorizable) portions of software. Putting effort into vectorizing your code on the Cray-1 simply demonstrated additional (often impressive) gains in performance.

During this innovative era in HPC dominated by CDC and Cray Research, Burroughs and the University of Illinois developed the *ILLIAC IV*, which introduced the concept of *parallel processing* as a viable model for dramatically increased performance. While the Illiac IV was primarily a research computer, it was ultimately delivered to NASA Ames in 1975 (partly due to the anti-defense politics of the era related to housing a computer at a public University), and demonstrated what would become a potent path to future HPC improvements through massively parallel processing.

In the early 1980's, another prolific hardware designer emerged within Cray Research named Steve Chen, who developed the *Cray XMP*. Taking a page from the lessons learned in the Illiac IV and other parallel designs, the XMP contained multiple CPUs. It was succeeded in the late 1980's with the *Cray YMP*. The YMP married the potent blend of fast serial computing (largely enabled by extremely fast memory), vectorized computing, and

modest parallelism with up to 8 CPUs. It also contained the first Unix-based operating system (UNICOS), which helped spawn an era of program portability in scientific computing.

Seymour Cray continued to design his own line of computers alongside the successful designs by Steve Chen. The Cray-2 was introduced in 1985, and was nicknamed “bubbles” due to the use of fluorinert for cooling, which was clearly visible through a transparent heat exchanger that added yet another element of visual appeal to the typical Cray designs of the time.

Continuing to push the boundaries of performance, Seymour Cray then started *Cray Computer Corporation* – another spinoff company focused on his own aggressive designs. Cray Computer started with development of the Cray 3, which innovated the use of Gallium Arsenide as a replacement for silicon. The company eventually killed the development of the Cray 3 due to cost overruns and delays. The Cray 4 likewise never made it to market due to Cray Computer going bankrupt in the early 1990’s.

Seymour Cray’s final attempt at redefining supercomputing was realized with the founding of *SRC Computers* in 1996, using the novel approach of Field Programmable Gate Arrays (FPGAs) to accelerate HPC. But the vision was forever hobbled when Seymour died in a tragic automobile accident just weeks after the founding of the company. SRC continues on to this day without Seymour’s leadership, but was never able to have the impact on the market that may have been realized had Cray been there to provide the vision.

Cray Research continued to define the supercomputing era throughout the remainder of the late 80’s and early 90’s with a series of upgrades. The Cray *C90*, and *T90* were introduced to the market, and continued to be the workhorse machines for major scientific laboratories in the United States. The Cray *J90* was also introduced in the early 90’s as a smaller, more affordable version of the Cray vector architecture that allowed for continued portability of applications at a price point more institutions (e.g. universities and businesses) could afford.

While the vector computing era was dominated by Cray Research and its competing lines of successful architectures, a number of other companies also found success during that era developing supercomputers based on vector processing. These include, but are not limited to: IBM, Convex, ETA (a spinoff of CDC), and Thinking Machines. The Japanese also developed a number of vector architectures through Fujitsu, NEC, and Hitachi – a foreshadowing of Japanese policy into the 1990’s and 2000’s that continued to rely on custom vector processors to achieve top performance.

While Cray and others helped push the boundaries of supercomputing through the 70’s, 80’s, and early 90’s – they could simply not compete with the increasingly cheap microprocessors that were finding broad market adoption. The invention of the first commercially successful personal computer by IBM in the mid 1980’s meant that computing was no longer the purview of government laboratories and large corporations. Computing began to become a part of the fabric of society at many levels, and the HPC industry in the U.S. likewise had to change in reaction to the marketplace realities.

Parallel computing and Commodity Processors

The rise of microprocessors in the late 1980’s and early 1990’s not surprisingly had significant impact not only on scientific computing, but on traditional mainframe computing as well. In what author Clayton Christensen later referred to as “*disruptive innovation*”, the emerging workstation and PC marketplace were putting powerful computing on the desktop at a fraction of the cost of mainframes per floating point operation. While Cray Research and a handful of others dominated HPC, mainframes such as the *IBM 360* series and *Digital Equipment Corporation’s PDP* and *VAX* were quickly becoming dinosaurs as a result of this disruption. A huge wave of RISC-based microprocessors began to flood the workstation marketplace. These include the SPARC (Sun), PA-RISC (HP), Alpha (DEC), PowerPC (IBM and Motorola), i860/960 (Intel), and the eventual x86 (Intel) line of CPUs that became the choice of IBM for use in the burgeoning Personal Computer market. This “attack of the killer micros” as it was deemed in a seminal talk at the *Supercomputing 1992* conference by Eugene Brooks from LLNL, led to the ultimate demise of specialized Supercomputers in the U.S. as microprocessors penetrated the lower end of the market through low-cost commodity CMOS processors, as well as the high-end market dominated by HPC vendors.

As initially demonstrated by the Illiac IV almost two decades prior and significant research efforts in the 1980’s spurred through DARPA investments – the path to higher performance computing clearly lay in the ability to exploit parallel computing. By pulling together collections of cheap microprocessors to work on solving smaller pieces of a larger problem concurrently, the limits of performance were seemingly unlimited.

Custom high-speed interconnects such as the hypercube, 3D torus, and the fat tree provided high bandwidth, low latency communication between processors in a style of programming referred to as *distributed memory message passing*. The downside was that this new style of parallel computing required programmers to completely rethink how to structure their algorithms and applications. The emergence of standardized message passing APIs and libraries such as PVM allowed for portability across

different platforms that in turn encouraged a flood of applications to move to this new model. In 1994, the programming model that would come to dominate scientific computing through to the current day was codified by a community effort called the *Message Passing Interface* (MPI). Rapid adoption was spurred on by a high quality free reference implementation developed at the DOE's Argonne National Laboratory called *MPICH* that ran on virtually every platform, including networked workstations.

In addition to message passing, *shared memory programming* also rode the commodity wave. Unlike message passing, where the programmer must explicitly manage the communication of data in a program between independent interconnected CPUs and memory, shared memory programming allowed the programmer to communicate data through a common shared memory address space using a technique called *threading* to manage multiple concurrent streams of execution. While conceptually simpler, shared memory programming was often tricky as the users had to carefully manage their programs to avoid difficult bugs due to race conditions, critical sections, and other programming pitfalls that were heretofore relegated to complex concurrent programs such as operating systems. Like message passing, shared memory programming found a standard in a compiler directives-based approach called *OpenMP*, which targeted platform portability and ease-of-use for common shared memory paradigms in scientific computing such as parallel loops, atomic operations, and barriers.

Now that many pieces of a supercomputer could be assembled using largely commodity processors and networks, the U.S. (and the world) saw an explosion of companies beginning to compete for a foothold in what came to be known as the *cluster* market. In addition to some of the big players like IBM, Cray Research, DEC, HP, SGI, and Intel, a number of smaller companies emerged with their own parallel designs, including *nCube*, *MasPar*, *Alliant*, *Multiflow*, *Kendall Square*, *BBN*, *Sequent*, *Convex*, *Encore*, *Meiko*, *Supertek*, *Floating Point Systems*, *Tera*, and surely others (most of which have since died off, or merged with larger companies).

In addition to these custom-built systems, the emergence of an open-source software stack built around the burgeoning Linux operating system spawned the *Beowulf* revolution, where just about anyone with access to a few PC's could experiment with parallel distributed-memory computing for the cost of the off-the-shelf hardware and enough tenacity to build and debug the cluster. This was a particular boon to universities and small businesses, and cemented the role of parallel computing as a preferred path to performance across the computing spectrum.

Current Day HPC

Today, the HPC industry is ensconced in the midst of yet another major shift in computing paradigms. This was brought about starting in the late 2000's when improvements in microprocessor performance could no longer ride solely on the continued increases in clock speeds due to the excessive heat such designs would incur. Initially, this played out largely through the introduction of multi-core processors, where a single CPU (or socket) consisted of multiple independent cores sharing a common main memory.

But in the never-ending quest for higher performance, a number of radical changes are taking place that is threatening to upend the comfortable combination of message passing and shared memory threading. Perhaps the most obvious incarnation of this is the introduction of accelerators, personified most clearly in the use of *General Purpose Graphics Processing Units* (GP-GPUs) spearheaded by companies like nVIDIA and AMD. While GPUs were once relegated to the task of accelerating 3D graphics for realistic gaming and high-end scientific visualization, companies are quietly adding features and programming environments that make them more appropriate for general purpose scientific computing.

In addition to GPUs, Intel introduced their answer to acceleration through the Xeon Phi line of co-processors (also known as the "MIC", or "*Many Integrated Core*") in 2010, which packs a large number of simpler cores on a single card optimized for HPC-style applications.

But as any HPC developer knows, increased performance relies not just on the speed of the CPU, but also the memory that feeds instructions and data to the CPU. For years, the industry has been worried about the impending *memory wall*, brought on by the increasing imbalance between processor and memory speeds that threaten to starve the processor by not being able to feed it data fast enough to process at speeds it is designed to operate at.

The computing industry has risen to the challenge through the development of 3D *stacked memory*, which promises to overcome the bandwidth restrictions of standard DRAM by stacking memory chips in the 3rd dimension. This provides a larger number of paths in and out of memory, versus the limits imposed by that of a 2D chip where pins are limited to the relatively small surface around the edge of the memory package. While stacked memory will likely help delay the issues of the memory wall for at least another few generations of processors, it is generally too expensive to use as the sole main memory in a compute node while still providing sufficient memory capacity for complex problems. Thus, the programmer will be faced with another type of heterogeneity by needing to manage the flow of data between the high-capacity DRAM and the high-bandwidth stacked

memory. This complex memory model will likely be further exacerbated as new non-volatile RAM technologies are introduced (such as *Phase Change Memory* (PCM), *Memristor*, and *Spin-Transfer Torque RAM* (STT-RAM)), which promise to provide much greater memory capacities, but at lower performance compared to stacked memory. Programmers will soon be faced with a daunting memory hierarchy that includes registers, cache (multiple levels), stacked memory, standard DRAM, NVRAM, burst buffers, and traditional “spinning disk” storage.

Finally, the HPC industry is taking a page from the *big data* community by inverting the relationship between memory and the CPU in the programming model. The *Map-Reduce* programming model pioneered by Google and made popular through the open source *Hadoop* framework took the novel approach of moving the computation to where the data lives versus bringing data over a bus or network to the CPU performing the task. In HPC, this is likely to be realized through a model called *processing-in-memory* (PIM), where a separate (perhaps much simpler) CPU is placed very close to the memory it operates on. This allows certain memory-intensive algorithms to be off-loaded from the main CPU, providing yet another level of parallelism (and heterogeneity) to emerging systems. While PIM is not a new concept in computing, it is re-emerging as yet another approach to help defeat the memory wall.

Big Data and “traditional” HPC are also seeing increasing signs of convergence. While initially data-centric systems focused on hardware typically eschewed by HPC (e.g. slow disks, commodity Ethernet, virtualization), many of the needs in data science are beginning to take advantage of HPC designs, and likewise are influencing the directions the HPC must head (e.g. larger memory and storage, and PIM). As big data matures from simply finding a needle in the haystack (search) to rapidly finding patterns in large data sets to provide unforeseen insight, HPC innovations such as accelerators and high-speed networks will influence the development of those systems. These data analytics techniques will likewise increasingly become part of the overall workflow of scientific simulation, especially as fields like *uncertainty quantification* (UQ) and *in-situ visualization* mature. Many in the field are seeing a natural convergence in HPC and Big Data coming, and the opportunities to redefine HPC to include this domain present exciting and important opportunities.

The Exascale Challenge

Many of the innovations described in the previous section would probably not be so urgent to solve for HPC if it weren’t for the fact that extremely large systems consisting of thousands or tens of thousands of nodes are needed to continue reaping ever-increased aggregate

performance. Just as the first Petaflop computers were beginning to reach the HPC market, DARPA released a paper (Kogge, 2008) outlining the challenges that would arise as the next factor of 1000x performance was attempted – known as an *Exaflop* (10^{18} floating point operations per second). Likewise, massive data requirements were going to require on the order of an *Exabyte* (10^{18} bytes) of storage, and the term *Exascale* refers to machines that reach either or both of those 10^{18} “Exa” goals. DARPA concluded that without radical changes to the design of large-scale systems (and thus the components on which they are built), the power requirements for an exascale system would reach into the 100’s of Megawatts. With the average cost of electricity in the U.S. hovering around \$1M/year/MW, the cost of keeping an exascale system powered up over a nominal 5 year lifetime would likely dwarf the capital cost of the hardware to assemble the system.

In addition, the sheer size and scale of an exascale system combined with potentially less-reliable components running at near-threshold voltages in order to save energy, will mean that *resilience* and *reliability* will increasingly become a concern for the application developer as well as the system designer. While standard checkpoint-restart techniques will likely continue to be utilized, applications and run-time systems will need to be increasingly fault-tolerant, either through finer-grained rollback to a known good state, or by algorithms that are naturally resilient to the occasional soft error.

These precise issues of power and resilience are at the heart of the exascale challenge, and is what’s driving ever-more creative ways to squeeze out additional performance-per-watt from CPUs, memory subsystems, network interconnects, and file systems.

The DARPA report was followed by a number of US DOE studies extolling the challenges of exascale computing, as well as the potential game-changing insights into a number of domains and missions that would occur if this level of computing power could be successfully harnessed. In addition to the hardware challenges of developing a usable system within affordable power budgets, the sheer magnitude of parallelism that must be exploited in these systems means that the design of algorithms and software must likewise again be reimaged and reinvented.

In the U.S., the DOE has been actively pursuing the formation of an *exascale initiative* for the past 5 years, and while there has been limited success in directing research budgets within the DOE’s Office of Science and NNSA to reach this goal, the initiative is yet to be approved and initial vendor partnerships have largely had to come through redirection of funds from other valuable research and development efforts. The original goal of

fielding an exascale system in 2018 within a 20MW power budget now is seen as highly improbable, and current proposed U.S. roadmaps are striving to achieve that goal in 2022 or 2023. Even then, the 20MW goal may prove to be too difficult to achieve on a system that is useful for general-purpose computing.

Meanwhile, other nations and the EU have not taken their foot off the accelerator, and perhaps seeing an opening in the U.S. laggard approach thus far, are racing to reach that goal well ahead of the U.S. Most HPC industry watchers agree that an exascale system by 2018 is technically within reach, but barring some sort of secret as-yet-unannounced technology, will likely have to brush aside concerns about power, energy, and productivity.

The race is on.

HPC LEADERSHIP: WHAT DOES THAT MEAN?

How exactly does a party interested in being considered a leader in HPC measure its progress? Part of the problem we have is that there is no satisfactory metric that captures what it means to be a global or national leader in HPC. While individual sites and installations have a standardized method in the *Top500* list, measuring true leadership is a more difficult metric.

Institutional Leadership

Top500 List

For the past few decades, the *Top500 Supercomputer Sites* list has almost exclusively filled the roll of defining HPC leadership. This site ranks the worlds' fastest 500 computers as defined by the *High Performance Linpack* (HPL) benchmark. While submission to the list is voluntary, it has provided perhaps the best and only platform for accurately measuring HPC trends over decades: the rise and fall of various microprocessors, vendors, interconnects, and countries.

The HPL benchmark is based on very dense matrix multiplication operations that in turn make extraordinarily efficient use of memory subsystems and vector/SIMD units relative to real applications of interest to most HPC users. It is not uncommon for HPL to report a floating-point operations per second (FLOPS) rate that is 75-85% of the peak FLOPS of the system. For many real applications, 5-10% peak performance is considered good, and anything approaching 40-50% is considered extremely fortunate. As a result, HPL has been increasingly criticized as a benchmark – for at best not being representative, and at worst driving vendors toward system designs that will benefit their Top500 ranking but not necessarily impact real application performance.

However, there is a strong correlation between how well a system performs in the Top500 and how well a general application is likely to perform. Relative performance as measured by HPL is likely a good enough indicator of whether other benchmarks will see similar relative performance, however we don't have hard historical data (yet) to prove or disprove that point. It is the case that most machines in the U.S. are procured based on other benchmark metrics besides HPL performance, and HPL is only run after the fact to generate and report a Top500 number for consideration in the ranking.

Achieving a ranking within the Top 10 of the Top500 certainly has its benefits, as it tends to get widely reported in the press, is anxiously followed in HPC circles, makes funding agencies happy, scores political points, and provides bragging rights to the vendor and site that partnered to deploy the system. While other benchmarks such as the *Green500* have emerged to attempt to measure aspects such as power efficiency of a system, and the *Graph500* for measuring data-centric performance of graph traversal class problems, the Top500 remains as the standard-bearer for declaring, "we're #1".

HPC Challenge Benchmarks

Partially to address the criticisms of the HPL benchmark being lopsided in its focus on floating point performance, the *HPC Challenge Benchmark* was developed around 2004 as a suite of seven benchmarks which collectively measured different aspects of the system, including memory and interconnect. While each individual component of the benchmark was narrowly focused, collectively the suite would more accurately predict system performance on real applications that stressed various aspects of the hardware. For a while in the late 2000's, HPCC was being suggested and promoted as a potential replacement (or at least an equal complement) for HPL in the Top500, but despite its benefits, it never received the momentum required to gain significant mindshare. Despite that, HPCC continues to be developed by the University of Tennessee, is still widely used, and continues to have its own award category at the annual SC conferences.

HPCG – An Emerging New Metric for the Top500

The maintainers of the Top500, with the support of the DOE, are again addressing this primary concern of HPL not being representative by proposing another change to the benchmark used. Instead of a dense matrix solve, the emerging HPCG benchmark uses a conjugate gradient-based matrix solver that tends to primarily stress memory bandwidth – a more common bottleneck for realized performance in many applications. While a single benchmark may never accurately capture the breadth of concerns that represent HPC applications that a suite of benchmarks like the HPCC might, a big benefit of HPCG is that it maintains all of the hallmarks that have made the

Top500 under HPL successful - namely, being easy to configure and run, small enough to optimize and take advantage of tuned math libraries, and reports a single metric.

As of early 2014, the HPCG benchmark was being introduced to a community of experts for feedback, and is undergoing continued tweaks before being formally announced as a potential complement to HPL. It is planned that for at least the foreseeable future, both HPL and HPCG results will be run in tandem. But since an unprecedented 20+ year history of exists based on HPL, it remains to be seen how quickly the community will adopt this new metric.

While the *Top500*, with its advantages and pitfalls is still the most widely recorded metric of institutional leadership, the community is beginning to make necessary changes to improve metrics.

National Leadership - Measuring the Intangibles

The country with the #1 machine or the most machines in the Top 10 of the Top500 is generally accepted as a first-order indicator of that country's relative strength in HPC. Competition for that #1 slot is intense, and a shift in the top spot between countries can spark some angst in the country being displaced, as happened in 2002 when Japan displaced the U.S. with the Fujitsu *Earth Simulator*.

However, the focus on measuring the benchmark performance of machines still falls well short of a satisfactory metric for understanding national leadership in HPC. While developing and deploying fast machines is the current best method available, it is lacking the ability to measure the far more important yet elusive metrics that define real HPC leadership. These include, but are not limited to:

- Breadth of HPC applications running at full scale
- Development of the supporting software stack
- Number of scientific discoveries enabled by HPC
- Return on investment (e.g. money saved by using simulation in lieu of traditional experiment)
- Academic achievements such as radically new scalable algorithms
- Adoption by domestic industry
- Ability to field large systems in lieu of embargoes
- Degree of innovation in tackling hardware challenges
- And of course – hard performance metrics like the current Top500 rankings.

Gathering metrics on the above would certainly allow for a more comprehensive global view of HPC leadership, but would require an effort on the scale of measuring the quality of American Universities in the popular and respected annual *U.S. News and World Report* rankings.

Even if the effort of that scale were put into place, it would be subject to missing data (either inadvertent or due to underreporting), likely be deemed as too “unscientific” by an otherwise scientific-minded community, and could potentially be subject to accusations of bias.

The market research firm International Data Corporation (IDC) collects and analyzes much of this type of data, which is then presented in summary form at major HPC conferences (including their own industry-focused HPC conference, the *HPC User Forum*). But the IDC analysis is more focused on spotting global trends and spending patterns versus attempting to create a concrete ranking.

SCIENCE AND COMPUTING POLICY IN THE U.S.

The seven-decade story told in the first portion of this paper was viewed mostly through the lens of technological advances of the supercomputer systems themselves, with government involvement coming primarily in the early years when it was the main customer, and the later years when it continued to advance the state of the art through agency-specific (primarily DOE and DoD) large scale procurements. But what was the role of the U.S. Government in shaping these changes, and would they have happened anyway? Would the DOE labs have been able to stay at or near the top of the global HPC race by procuring ever-larger machines, or were there other factors at play in the halls of the executive and legislative branches of government that helped as well?

To answer these questions, it is first helpful to navigate the sometimes-esoteric way in which policy is made in the United States. Through the miracle of democracy on which this country is founded, policymaking involves a delicate balance between the executive and legislative branches of government, the various agencies tasked with implementing HPC to deliver on their missions, and the U.S. computing vendors who provide the technology.

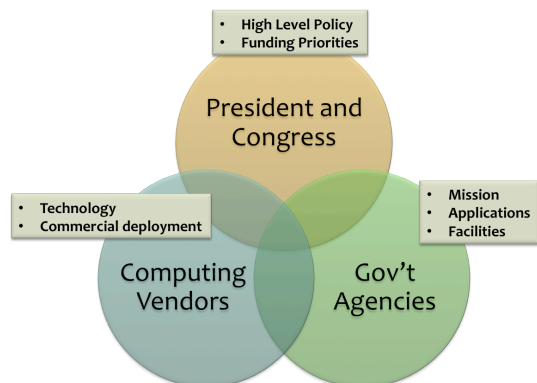


Figure 2 - HPC Policy in the U.S. involves politicians, gov't agencies, and private sector technology providers working together

The Executive Branch

The annual U.S. budget cycle begins the year before when the President releases his version of the budget for the following fiscal year. While this is rarely the budget that ultimately gets passed by Congress, it serves an important role in helping to define the priorities of the *Executive Office of the President* (EOP), and at least in theory is used as a starting point for Congress to begin the task of appropriating real dollars to various discretionary programs and initiatives.

In the EOP, HPC policy has largely been reflected in the budget through recommendations made by the *Office of Science and Technology Policy* (OSTP) and the *Office of Management and Budget* (OMB). Each year these two organizations issue a Memo of Science & Technology Priorities to the various agencies highlighting science policy shifts and/or priorities. While there is no budget attached to these policy recommendations, they are still an important reminder to the various agencies what the EOP expects them to be focusing on.

While the EOP and OSTP have largely been silent on advanced computing in the past decade, in March of 2012 the Obama administration (through the OSTP) initiated a “Big Data” Initiative, which directed \$200M to be spent across multiple agencies to “*improve our ability to extract knowledge and insights from large and complex collections of digital data*”. While big data is an important element of a comprehensive HPC policy, this policy lacked any ties to the challenges and benefits of advancing traditional numerical modeling and simulation, which in turn did not help DOE’s case for launching their proposed exascale initiative. This is expected to change in 2014, as OSTP has initiated a Fast Track Action Committee (FTAC) to establish the *National Strategic Computing Initiative*, or NSCI. The goal of the NSCI is to “*create a cohesive, multi-agency strategic vision and investment strategy that assures the United States sustains or extends its historical lead and strategic advantage in High Performance Computing (HPC) technology for national security, economic prosperity and scientific discovery*”. This effort was commissioned by the Committee on Homeland and National Security, in the National Science and Technology Council within the OSTP, and is expected to comprehensively address both big data and exascale, as well as the future challenges of “beyond CMOS” technologies as existing silicon-based processors reach their physical limits of manufacturing technology sometime in the next decade.

The Legislative Branch

The U.S. Congress is the branch of government that is ultimately responsible for passing the annual budget, and thus has the ability to put real money behind national

priorities. The *House Committee on Science, Space, and Technology* currently chaired by Lamar Smith (R-TX) in the 113th Congress, is where most of the House of Representatives discussions related to HPC occur. The *Senate Committee on Commerce, Science, and Transportation* and the *Senate Committee on Energy and Natural Resources* also hold occasional hearings, but to a lesser extent than the House Committee.

Members of the Senate have expressed strong support for the DOE’s exascale initiative in 2011 and again in 2013, through a bi-partisan letter to the president. The Senate under the leadership of Lamar Alexander (R-TN) submitted senate bill S.733 in April of 2013 titled the *ExaSCALE Computing Leadership Act of 2013*. However, this bill never made it out committee for broader discussion and debate. Likewise the House of Representatives wrote a bill which directed DOE to “*develop and carry out a plan to develop and incorporate exascale computing ... into the DOE stockpile stewardship program*” in the *2014 National Defense Authorization Act* (HR 3304, Section 3129) which became law over the Christmas Holiday of 2013.

ADVISORY GROUPS

Both the executive and legislative branches rely on expert advisory groups to help distill complex science and technology issues so either OSTP or Congress can act upon them in an informed manner. These groups are designed not to lobby for specific legislation, but to instead provide expert advice to help craft effective policy, which in turn can become legislation with appropriated funds.

PCAST

The OSTP relies on advice and input from several different groups. One of the most transparent and influential is the *President’s Council of Advisors on Science and Technology*, or PCAST. With its roots going back to the Roosevelt Administration in 1933, PCAST is made up of approximately 20 non-government experts in science and engineering whose charter is to advise the president (and thus the OSTP) on policy matters in science and technology. PCAST reports are regularly published for anyone to read. PCAST has not yet expressed explicit support for a well-defined exascale or next-generation computing initiative.

JASON

JASON is an independent scientific advisory group established in 1960 that provides reviews and recommendations on a myriad of topics to the U.S. government – largely focused on matters of national defense. The JASONS have been chartered several times to study matters related to the use of HPC, mostly related to national security missions housed in Department of

Energy's *National Nuclear Security Administration* (NNSA). In 2003, JASON released a report titled *Requirements for ASCI* (Schwitters, 2003), which reviewed the NNSA's *Accelerated Strategic Computing Initiative* (discussed later in this paper). It concluded that ASCI (and thus HPC) had become "essential to the Stockpile Stewardship mission". In 2013, JASON released *The Challenges of Exascale Computing* (McMorrow, 2013), with conclusions that confirmed the extreme challenges that lay ahead, and recommended continued strong government investment. But the 2013 JASON report also challenged the SSP mission need to deploy exascale class systems by 2020 (a stated goal at the time of the review) while the software challenges were seemingly even larger – thus calling for a moderated timeline and cautious approach to deploying the hardware before the applications had solved some of their fundamental issues.

NITRD

The *Networking and Information Technology Research and Development* (NITRD) Program was borne out of the *High Performance Computing Act of 1991* (discussed below), and is designed to stimulate coordination across many agencies involved in computing and information technology, with a focus on high-end computing. NITRD is subdivided into eight *Program Component Areas* (PCAs) that are coordinated by various *Coordinating Groups* (CGs), *Senior Steering Groups* (SSGs), and *Interagency Working Groups* (IWGs) who report on budgets, coordinate planning, assist with policy recommendations, and collaborate in areas of common interest. In particular, there are two PCA's specifically targeted at HPC, the *High End Computing Infrastructure and Applications*, and *High End Computing Research and Development*.

NITRD releases an annual supplement to the President's budget request commonly referred to as the "Blue Book". The most recent supplement released for the 2015 budget cycle contains a breakdown of activities across the agencies represented in NITRD related to the HEC and other PCAs.

NITRD also publishes strategic planning documents, with the most recent update in July of 2012. This most recent strategic plan discusses the challenges of exascale, in particular the research needs related to resilience and energy.

In 2003 the *HEC Revitalization Task Force*, under the aegis and direction of NITRD, held a workshop and subsequently released a report titled *The Roadmap for the Revitalization of High End Computing* (Reed, 2003) that had impact on helping multiple agencies (DOE NNSA, DOE SC, NASA, and NSF) obtain increases in their base budgets for HPC. This report also likely had positive

influence on the passing of the *DOE High-End Computing Revitalization Act of 2004* (discussed below).

SEAB Task Force on Next Generation HPC

In late 2013 continuing into 2014, the DOE Secretary of Energy Ernest Moniz tasked the *Secretary of Energy Advisory Board* (SEAB) with the formation of a *Task Force on Next Generation High Performance Computing*. This task force consists of high-level non-government experts in HPC across academia, industry, computing vendors, and select (non-DOE) agencies. They focused on the justification for a DOE Exascale program as related to mission need, fundamental research opportunities, and broader societal impacts that such an initiative will enable. A draft of the report was formally submitted to the Secretary of Energy at the June 2014 SEAB meeting, and subsequently made publicly available in August 2014 (SEAB, 2014).

Council on Competitiveness

The Council on Competitiveness (CoC) is a non-profit, non-partisan, non-government organization created during the Reagan administration in 1986, and is composed of membership by corporate CEOs, University presidents, labor leaders, and national laboratory directors. The Council uses these thought leaders to craft recommendations to U.S. policy makers on a variety of policy matters related to driving U.S. competitiveness and productivity.

The CoC is unique among organizations of this type in that for the past decade, it has identified High Performance Computing as a primary driver for competitiveness in the U.S., and has worked diligently to help spread the message of the importance HPC through policy recommendations, case studies, market research, and the formation of public-private partnerships. The Council's *HPC Advisory Committee*, consisting of HPC experts from their membership, has been actively debating HPC policy since 2004 and has released a number of case studies and reports. In 2011, the Council led the formation of the *National Digital Engineering and Manufacturing Consortium* (NDEMC) with a pilot program in the U.S. Midwest focused on spurring the use of HPC in small and medium sized U.S. manufacturers. This public-private partnership provided these smaller companies that are critical to the manufacturing supply chain access to computers and software at several large academic centers (NCSA and University of Illinois, Ohio Supercomputer Center, and Purdue University), along with training and software.

Office of Technology Assessment

It is worth noting that in 1972, Congress developed the *Office of Technology Assessment* (OTA) that was chartered to advise Congress on all matters of science and technology. The office was closed in 1995 as part of an

effort to shrink government under the Republican-led “Contract with America” despite its relatively lean \$22M annual budget. Calls to reinstate the OTA have come from various members of congress, leading scientists, and the occasional presidential candidate. Without a single legislative entity to study and advise for advances in science and technology (including HPC), it is often hard to find a unified voice for advocacy in Congress.

U.S. GOV’T AGENCIES SPONSORING HPC ADVANCEMENT

There is not one U.S. government agency solely responsible for the advancement of science and technology in the United States. One can debate whether this is a good or bad thing, but it is worth noting that other countries with technology leadership or leadership ambitions have the equivalent to a non-existent “U.S. Department of Science and Technology”. The United Kingdom has the *Ministry of Technology*, India has the a *Ministry of Communications and Information Technology*, South Korea has the *Ministry of Science*, and China has the *Ministry of Science and Technology* (to name a few).

According to the best public data available from NITRD and presidential budget requests, the three agencies consisting of the Department of Energy, Department of Defense, and National Science Foundation collectively form the majority of HPC expenditures in U.S. government.

Department of Energy

The DOE consists of the *Office of Science* (SC) and the *National Nuclear Security Administration* (NNSA) – both of which are continuously driving high end scientific computing. Combined they make up close to half of the reported budgets for HPC deployment and development in the U.S. The Office of Science, led by the *Leadership Computing Facilities* at Argonne and *Oak Ridge National Laboratories* provide unclassified computing to a variety of customers within DOE, academia, industry, and other agencies. DOE Applied Programs are supported by the *National Energy Research Scientific Computing* (NERSC) Center at *Lawrence Berkeley National Laboratory*, as well as HPC centers at the *National Renewable Energy Laboratory* (NREL), *Pacific Northwest National Laboratory* (PNNL), and the *National Energy Technology Laboratory* (NETL). Classified computing centers at *Lawrence Livermore National Laboratory* (LLNL) and the *Advanced Computing at Extreme Scale* (ACES) Consortium consisting of *Los Alamos* and *Sandia National Laboratories* in New Mexico lead the NNSA HPC efforts.

Accelerated Strategic Computing Initiative (1995)

In 1995, the DOE launched the *Accelerated Strategic Computing Initiative* (ASCI) Program. ASCI was conceived at what was in retrospect a perfect storm of mission need, policy momentum, and technological breakthrough – and helped cement the U.S. lead in HPC for the decade that followed and beyond.

In the early 1990’s, the U.S. ceased underground testing of nuclear weapons. Up until that point, simulation was largely used to assist the weapon designer, but it was not until an underground nuclear test was performed that true confidence in the performance or safety of a design would be trusted. The cessation of testing, the aging (and thus retirement) of experienced designers, and the aging U.S. nuclear stockpile provided exactly the sort of time-critical mission need required to change business-as-usual. With the advent of distributed memory parallel computing just starting to take root at this time, this confluence of events launched a new program of *science-based stockpile stewardship*, rooted in the simulation advances promised by the ASCI program along with sophisticated new experimental facilities.

As part of the ASCI Program, a series of *Academic Alliances* were formed starting in 1997 that aimed to help prepare the workforce for NNSA to address their challenges in the Stockpile Stewardship Program. Currently in its third incarnation, five-year awards are given to Universities to establish open and unclassified centers focused on challenging problems in physics and the application of HPC to solve those problems. The centers are aimed to model the type of large multi-disciplinary projects typical in the NNSA, and serve as a “training ground” for the type of mission-driven research critical to ASCI’s continued success. Many graduates of these universities go on to work at NNSA laboratories.

While the large ASC computers at the NNSA are used primarily by the NNSA Laboratories to address mission needs in stockpile stewardship, the DOE SC Leadership Computing Facilities are open unclassified systems largely available through grant application process to researchers (discussed in more detail below).

Department of Defense

The Defense Advanced Research Projects Agency (DARPA) has been a primary driver of HPC research in the DoD. Most notably, DARPA has pushed the boundaries of computing technology through aggressive initiatives such as the *Strategic Computing Initiative*, the *High Productivity Computing Systems* (HPCS) Program, and the *Ubiquitous High Performance Computing* (UHPC) program. DARPA’s focus is largely on developing the foundational research in HPC.

The DoD *HPC Modernization Office* supports five *DoD Supercomputing Resource Centers* (DSRCs) housed at various military bases for the Army, Navy, and Air Force. These large centers provide centralized supercomputing services to all branches of the Military, and while not typically of the size and scale of DOE Supercomputing centers, represent a significant amount of compute power for both unclassified and classified computing needs at the DoD.

DARPA Strategic Computing (SC) Initiative (1983)

In the early 1980's, DARPA initiated an aggressive program targeting a goal of realizing the dreams of the Artificial Intelligence (AI) community – an intelligent machine that could “think” like a human (Roland, 2002). Tying together research in chip design, computer architectures, and software, the program ended up having broad impact on the development of parallel computing technology. The computer architectures program element quickly realized the need for parallelism (in an era when vector computers dominated HPC) to reach their goals of machine intelligence, and developed and matured some of the early concepts behind parallel computing. The first architecture project authorized in 1983 was the *BBN Butterfly*, a shared-memory design built on Motorola CPUs and a novel interconnect that gave processors access to each other's memory banks. In an attempt to ensure architectural diversity, noted MIT architect Danny Hillis established funding through DARPA to realize his Connection Machine design – an early example of massive parallelism across very simple processors, or “cells”. Hillis went on to form *Thinking Machines*, one of the early innovators in parallel computing when the technology finally found firm roots and a commercial market almost a decade later.

The Strategic Computing Initiative ultimately spent \$1 billion dollars without ever reaching the ultimate goal of a “thinking machine”, but far from being a failure – the program simply redefined success as sowing the seeds for the explosion of parallel computing in the 90's.

DARPA High Productivity Computing System (HPCS) Program (2002)

In the early 2000's, DARPA initiated the High Productivity Computing System (HPCS) Program as an approach to ensure that U.S. HPC innovations didn't lapse by resting on the laurels of the robust linux-based cluster market at the time, while also maintaining a focus on designing systems that addressed the overall productivity of the developer and user. The goals of HPCS focused on four areas: Performance, Programmability, Portability, and Robustness.

The program was structured with a phased approach over approximately 8 years, where the beginning of each phase involved a downselect of technologies that would move

forward into the next phase, with an end goal of deploying multiple marketable systems in the 2011 timeframe. Vendors were encouraged to think about both hardware and software in the development of their systems, including novel new programming languages co-designed with the hardware platform that would inherently understand concepts of parallelism in the core language, and aim to dramatically improve programmer productivity.

The HPCS Program ended in 2011 with the successful development of two major systems: The *IBM PERCS* (Power7-based) with the *X10* programming language, and the *Cray Cascade* (Intel Xeon-based) with the *Chapel* programming language. IBM went on to market the PERCS system as the *Power 775*, and Cray marketed the Cascade system as the *XC30*. The X10 and Chapel programming languages likewise continue to be developed but are still years away from providing a robust and portable choice for large-scale application development.

National Science Foundation

The NSF also supports a number of supercomputing centers aimed at unclassified and academic research in advanced computing infrastructure and scientific discovery. Like the DOE, the NSF awards large blocks of time on supercomputers through the *XSEDE* program – a collection of mid-to-large sized computing centers largely housed at U.S. Universities and managed as a single virtual system providing HPC and large data services to the research community.

NSF likewise supports the *Blue Waters* system at the *National Center for Supercomputing Applications* (NCSA) at the *University of Illinois*. Currently the largest system housed at any university, grants for time on Blue Waters are provided through the NSF's *Petascale Computing Resource Allocations* (PRAC).

Like DOE, the NSF has likewise also understood, maintained, and emphasized the importance of training a workforce capable of making effective use of HPC to a broader audience.

Other Agencies

Rounding out the HPC landscape in the U.S. government are a number of smaller installations at the *National Institute for Health* (NIH), *National Aeronautics and Space Administration* (NASA), the *National Institute for Standards Technology* (NIST), the *National Oceanic and Atmospheric Administration* (NOAA), and the *National Center for Atmospheric Research* (NCAR).

COMPUTING POLICY INITIATIVES

Up until the early 1990's, computing policy in the U.S. was largely defined by the government use of large computers, and the development of large scale networking projects such as *ARPANET* and *NSFnet*. These government-supported networking innovations allowed researchers to access the power of computers through wide-area computer networks, and thus greatly expanded the accessibility of these resources to an ever-broadening audience.

Starting in the mid-1980's, global competition in computing began to take on increased urgency, and the U.S. Government reacted by shifting their role to include more policy-driven approaches. What follows is a sampling of some high level U.S. Policy initiatives that have had a positive impact on the U.S. leadership role in HPC.

SEMATECH (1987)

Up until the early 1980's the U.S. Semiconductor industry largely stood alone and led the world in semiconductor manufacturing. By the mid-1980's, competition from Japan and others caused the U.S. to lose that leadership, and the industry was in grave danger of completely shifting design and manufacturing offshore – a move that would be difficult if not impossible to reverse given the tremendous momentum in the growth of computing at the time. In reaction, the U.S. government spurred the creation of the *Semiconductor Manufacturing Technology (SEMATECH)* Consortium in 1987, aimed at coordinating R&D efforts by chipmakers, suppliers, research institutions, universities, and government agencies. The DoD and DARPA subsidized funding of this public-private partnership to the tune of \$500M over the first five years, and SEMATECH is largely credited with stemming the outflow of manufacturing and saving the U.S. semiconductor industry from potential obsolescence.

The High Performance Computing Act of 1991

ARPANET and NSFnet laid the groundwork for what would become the Internet as we know it today. But it was the *High Performance Computing Act (HPCA) of 1991* that helped kick the Internet into high gear. Sponsored by then Senator Al Gore, and partially in reaction to an increasingly loud groundswell of cries for action from the private sector (Crawford, 1991), the *High Performance Computing Act of 1991* (commonly referred to as the HPCA, or “Gore Bill”) left a huge legacy on the development of the Internet. Some of its more widely publicized success stories include the development of the *Mosaic* web browser at the NCSA and the *National Information Infrastructure* (NII), the combination of technology and policy that helped spark the Internet revolution.

While the focus of the HPCA was largely around networking and the Internet, the bill outlined a broad set of goals and important policy recommendations for high end computing. From a summary of the bill:

- 1) Setting goals and priorities for Federal high-performance computing research, development, and networking;
- 2) Providing for interagency coordination
- 3) Providing for oversight of the operation and evolution of the National Research and Education Network provided for in this Act
- 4) Improving software
- 5) Acceleration of high-performance computer system development
- 6) Technical support and research and development of software and hardware needed to address fundamental problems in science and engineering (Grand Challenges)
- 7) Educating undergraduate and graduate students
- 8) Providing for security.

Specifically, the HPCA directed the *DOE*, *NSF*, *NASA*, *NIST*, *NOAA*, and even the *Environmental Protection Agency* (EPA) and *Department of Education* to direct research toward advancing High Performance Computing. Likewise, NITRD was formed at the direction of the HPCA.

The Department of Energy High-End Computing Revitalization Act of 2004

Partially in reaction to the shockwave sent through the U.S. with the deployment of the Japanese *Earth Simulator* system in 2002, the 108th Congress passed Public Law 108-423 titled the *Department of Energy High-End Computing Revitalization Act of 2004*. This law directed the DOE to establish what became the *Leadership Computing Facilities (LCF)* located at Oak Ridge and Argonne National Laboratories.

These new user facilities were designed to provide high-end computing access and technical support on a competitive, merit-based review process to researchers in U.S. industry, institutions of higher education, national laboratories and other Federal agencies. In addition to providing tens of millions of CPU-hours annually per project, a focus of these facilities is to broaden the base of researchers able to effectively use HPC. Access to time on these LCF machines is broken down through three primary categories: The *INCITE* program (60%), *ASCR Leadership Computing Challenges* (ALCC) (30%) and *Director's Discretionary* (10%). In all cases, the focus of the LCFs is on large capability systems, and since its inception, both Oak Ridge and Argonne have held spots at the top of the Top500, including the #1 spot by *Titan* in 2012.

INCITE (*Innovative and Novel Computational Impact on Theory and Experiment*) aims to deliver high-risk, high-impact science that can only be performed on extreme scale systems. Successful peer-reviewed proposals are awarded annually across a range of fields, including material science, energy research, and climate change. Uniquely, INCITE also offers industry researchers access to grants that typically amount to about 5-10% of the total INCITE allocation going to the private sector. This in turn helps U.S. industry evaluate the potential return-on-investment of potentially deploying its own systems or purchasing time elsewhere for proprietary work.

Similarly, ALCC grants are awarded large blocks of compute time on either the *Leadership Computing Facility* machines or on the *NERSC* system housed at Lawrence Berkeley National Laboratory. Awards are also on a peer-reviewed merit basis, and are focused on grand challenges specifically facing the DOE mission.

America COMPETES Act (2007, 2010)

In 2007, and reauthorized in 2010, the president signed into law the *America Creating Opportunities to Meaningfully Promote Excellence in Technology, Education, and Science* act, or more commonly known as simply the *America COMPETES Act*. While the act does not explicitly call out a focus on computing, it did get broad bi-partisan support for reinvigorating STEM education, and focused on revitalizing manufacturing in the United States. Manufacturing is widely seen as some of the “low hanging fruit” to introduce HPC modeling and simulation to the broader industrial sector.

The American Recovery and Reinvestment Act of 2009

In the wake of the Great Recession of 2008, then newly-elected President Barack Obama and the 111th U.S. Congress enacted the *American Recovery and Reinvestment Act of 2009* (ARRA) as one of the first major initiatives of the new administration. The Recovery Act pumped approximately \$800B into the U.S. economy across a huge swath of “shovel ready” projects – many around exploring clean energy solutions, beefing up transportation infrastructure, and tax breaks for struggling Americans.

However the ARRA did provide a temporary boost for HPC projects as well – mostly notably the \$69M *Magellan* project at Argonne and Lawrence Berkeley National Laboratories designed to explore the emerging Cloud Computing market as applied to scientific computing. ARRA also provided \$125M for numerous desktop computing upgrades across many agencies, and \$90M for support of graduate students, post-docs, and PhD scientists – some of which helped buoy HPC research in the Office of Science, NSF, and elsewhere.

DOE FastForward and DesignForward (2012/2013)

In 2012, the Department of Energy kick started the push to exascale through the FastForward program, a \$62M competitively based award to computing vendors to explore novel research in processor, memory, and storage technologies. Awards were made to IBM, Intel (2), NVIDIA, and AMD and in conjunction with the ASCR and ASC co-design centers launched in 2012, helped initiate the co-design model where applications developers and hardware vendors work side-by-side on developing technology suitable for usable exascale-class systems. In April 2014, DOE released an RFP for the *FastForward2* program designed as a follow-on program to continue developing novel memory and node architectures in advance of current industry roadmaps.

Likewise, DOE initiated the \$25.4M *DesignForward-Interconnect* Program in 2013, awarding contracts to a number of vendors to explore novel interconnect technologies. IBM, Cray, Intel, nVIDIA, and AMD were all awarded contracts, and along with FastForward are leveraging ongoing co-design activities within DOE to insure innovative network interconnects are engineered with real applications in mind.

U.S. HPC POLICY RECOMMENDATIONS

HPC Policy in the U.S. has been effective over the past several decades to help ensure a robust and growing HPC ecosystem. Despite the lack of a single organization tasked with advancing HPC Policy, one might argue that the existing mixture of democracy, government purchases driving innovation, world-class academic institutions, and a demand from industry have done a fair job of helping the U.S. sustain leadership.

But a great number of thought leaders in the field are rightly concerned that with the tumult in Washington, the challenges of exascale, the surge of global competition, and the fact that HPC is no longer an appealing high-margin market as it once was for HPC vendors – the lack of a coherent U.S. HPC Policy going forward combined with the lack of long-term committed funding will negatively impact the U.S. leadership position. Particularly if competing countries aggressively take up the mantle of achieving efficient and effective exascale class platforms and applications over the coming decade.

What follows is a high level set of recommendations for a unified, coherent, and sustainable U.S. HPC Policy.

- 1) **Shift the meaning of “leadership”** away from institutions topping the Top500 list to a more national perspective. Develop the metrics necessary for a new method of ranking, and work with the global HPC

community to establish a trusted ranking that will benefit all nations in their quest for HPC leadership.

- 2) **Provide sustained funding for a DOE-led Exascale Initiative** to reinvigorate innovation, spur public-private partnerships with vendors, and provide affordable Petascale computing to academia and industry. **Fund the development of foundational math, computer science, applications, and data science** as a critical element to success. **Leverage the growth of “big data”** to extend the sphere of what it means to do HPC beyond traditional numerical physics-based simulation, and allow HPC simulation to ride the commodity wave driven by the data science revolution currently under way.
- 3) **Work closely with industry to expand their access to leading government computers and the expertise in the national laboratories.** The goal must be to demonstrate the return-on-investment newcomers will realize from shifting to a simulation-based design cycle. Part of such a process would require a reexamination of current policies regulating access to government resources, as these currently add significant latency and complexities to initiating these partnerships.
- 4) **Develop an enduring HPC workforce** (both through comprehensive indigenous training, and more liberal immigration policy for skilled workers) who can deploy this knowledge in U.S. industry and R&D organizations. For example, this should involve an aggressive expansion of government-academic partnerships.
- 5) **Leverage advances in cloud and mobile computing** to provide easier access to large resources and applications for a broader base of users. This includes both easier (e.g. pay-per-use) access to large-scale computing through web-based services, as well as focusing on providing validated solutions through easier-to-use interfaces that abstract much of the underlying complexities inherent in simulation codes so they can be used on focused problem areas by non HPC experts.
- 6) **Think “beyond exascale”** to perform the basic research now that will define the post-CMOS era of computing for sustained leadership, and so we are not simply building technology that will require another complete rethink from scratch in the decade that follows.

While government policy makers cannot and should not “pick technology winners” to bet on, a broad over-arching policy framework with clear mileposts, goals, funding, and outreach will provide the necessary spark so that HPC

in the U.S. can leverage the creativity and innovation inherent in the United States.

CONCLUSION

The United States has been a leader in High Performance Computing since the inception of the first computer, through a potent combination of private sector innovation and government support in the form of both policy and mission-driven government agency procurements. HPC leadership is now ours to lose, and the competition is fierce and growing. We are perched on a slippery slope in 2014. If we do not act aggressively, we will find it either impossible or overly expensive to regain what could be lost. U.S. Policy makers have a shrinking window of opportunity to help provide a coherent and broad set of multi-agency recommendations along with sufficient funding to continue working with U.S. industry in the decadal process of bringing radically new innovations to the mass HPC market.

We cannot simply ride on historical precedent, but must be increasingly creative in our policy approaches. Industry, academia, and multiple agencies must all be strong partners in this quest, requiring novel public-private partnerships, a focus on workforce training, and a recognition by industry thought-leaders that the “trickle down” effects of U.S. leadership in HPC will provide long-term economic security.

ACKNOWLEDGMENTS AND DISCLAIMER

This work was performed through the support of the *University of California Institute on Global Conflict & Cooperation*, and under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

BIBLIOGRAPHY

- Alexander, L. (2011, 11). *Letter to the President on Supercomputing*. Retrieved from <http://www.energy.senate.gov/public/index.cfm/2011/11/letter-to-the-president-on-supercomputing>
- ASC Co-design Project. (n.d.). Retrieved from <https://asc.llnl.gov/codesign>
- ASCI. (n.d.). *ASCI Academic Strategic Alliances Program Statement*. Retrieved from https://asc.llnl.gov/alliances/alliances_archive/asap/statement.html
- ASCR. (n.d.). *ASCR Workshops and Conferences*. Retrieved from Advanced Scientific Computing Research, DOE Office of Science: <http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>
- ASCR. (n.d.). *SciDAC/ASCR Co-design Centers*. Retrieved from <http://science.energy.gov/ascr/research/scidac/co-design/>
- Carr, A. (2013, April). *From Punched Cards to Petaflops*. Retrieved from National Security Science: http://www.lanl.gov/newsroom/publications/national-security-science/2013-april/_assets/docs/punchcards-petaflops.pdf
- Christensen, C. (1997). *The Innovators Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Review Press.
- Council on Competitiveness. (n.d.). Retrieved from HPC Initiative Publications: <http://www.compete.org/about-us/initiatives/hpc/>
- Crawford, D. (1991). CEOs Unite to Influence U.S. Technology Policy. *Communications of the ACM*, 34 (6), 15-18.
- Denenberg, S. (1971). *A 10 page description of the ILLIAC IV*. Retrieved from <https://archive.org/details/10pagedescriptio00dene>
- Department of Energy High-end Computing Revitalization Act of 2004 (Public Law). (2004). Retrieved from <http://www.gpo.gov/fdsys/pkg/PLAW-108publ423/html/PLAW-108publ423.htm>
- DOE/ASCR. (n.d.). *ASCR Leadership Computing Challenge*. Retrieved from <http://science.energy.gov/ascr/facilities/alcc/>
- DOE/ASCR. (n.d.). *INCITE Program Description*. Retrieved from DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program: <http://www.doeleadershipcomputing.org>
- Dongarra, J. H. (2013). *Toward a New Metric for Ranking High Performance Computing Systems*. Sandia Report.
- Dongarra, J. (2007). The HPC Challenge Benchmark: A Candidate for Replacing LINPACK in the Top500? 2007 SPEC Benchmark Workshop.
- Executive Office of the President & National Science and Technology Council. (2012). *NITRD Program 2012 Strategic Plan*. Retrieved from http://www.nitrd.gov/pubs/strategic_plans/2012_NITRD_Strategic_Plan.pdf
- Geist, A. e. (1994). *PVM: Parallel Virtual Machine. A Users Guide and Tutorial for Networked Parallel Computing*. MIT Press.
- Goldstine, A. (1946). The Electronic Numerical Integrator and Computer (ENIAC). *Mathematical Tables and Other Aids to Computation*, 2, 97-110.
- Kogge, P. e. (2008). *Exascale Computing Study: Technology Challenges in Achieving Exascale Systems*. DARPA.
- LARC. (n.d.). Retrieved from Computer-history.info: <http://www.computer-history.info/Page4.dir/pages/LARC.dir/LARC.Cole.html>
- Larzelere, A. (2009). *Delivering Insight: The History of the Accelerated Strategic Computing Initiative*. LLNL Technical Report UCRL-TR-231286.
- Liu, N. C. (2012). On the Role of Burst Buffers in Leadership Class Storage Systems. *MSST/SNAPI*.
- McMorrow, D. (2013). *JASON Report JSR-12-310, Technical Challenges of Exascale Computing*. JASON.
- NITRD. (2014, May). *2015 Supplement to the President's Budget*. Retrieved from <http://www.nitrd.gov/pubs/2015supplement/FY2015NITRDSupplement>
- NITRD. (n.d.). *NITRD Supplements to the Presidents' Budget*. Retrieved from <http://www.nitrd.gov/Publications/SupplementsAll.aspx>
- NSF. (n.d.). *National Science Foundation Advisory Committee for Cyber infrastructure Task Force on High Performance Computing*. Retrieved from http://www.nsf.gov/cise/aci/taskforces/TaskForceReport_HPC.pdf
- OSTP. (2012, March 29). Obama Administration Unveils "Big Data" Initiative: Announces \$200M in New R&D Investments.
- Reed, D. (. (2003, June). *The Roadmap for the Revitalization of High-End Computing*. Retrieved from <http://archive.cra.org/reports/supercomputing.pdf>
- Roland, A. a. (2002). *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1993*. MIT Press.
- Schwitters, R. (2003, October). *JASONS Report JSR-03-330, Requirements for ASCI*. Retrieved from <http://www.fas.org/irp/agency/dod/jason/asci.pdf>
- SEAB. (2014, August 10). *Report of the Task Force on High Performance Computing of the Secretary of Energy Advisory Board*. Retrieved from <http://energy.gov/sites/prod/files/2014/08/f18/SEAB%20HPC%20Task%20Force%20Final%20Report%2008-10-2014.pdf>
- SEAB Task Force on Next Generation High Performance Computing. (n.d.). Retrieved from <http://energy.gov/seab/secretary-energy-advisory-board-seab-task-force-next-generation-high-performance-computing>
- Sterling, T. (2001). *Beowulf Cluster Computing with Linux. Text of the High Performance Computing Act of 1991*. (n.d.). Retrieved from http://www.nitrd.gov/Congressional/Laws/pl_102-194.aspx
- The Manhattan Project Heritage Preservation Association. (2004). *Manhattan Project History: Evolving from Calculators to Computers*. Retrieved from <http://www.mphpa.org/classic/HISTORY/H-06c18.htm>
- von Neumann, J. (1945). *First Draft of a Report on the EDVAC*. Retrieved from <https://archive.org/details/firstdraftofrepo00vonn>
- Wulf, M. (1995). Hitting the Memory Wall: Implications of the Obvious. *Computer Architecture News*, 23 (1), 20-24.